Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 01:22:21
REFERENCE NO: 220

Author Names & Affiliations

- Frank Wuerthwein - UCSD
- David Swanson - UNL

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Physics/Particle Physics, Computer Science

Title of Submission

Integrated Cyberinfrastructure for Open Science

Abstract (maximum ~200 words).

Science is a team sport that increasingly involves multiple institutions, all of which have made significant investments into shared campus-wide Cyberinfrastructure. The team's success depends on dynamic pooling of any resources its members have access to. This typically includes resources they or their institutions own "in house", allocations they have at national-scale facilities, and resources dynamically purchased from commercial clouds as needed.

We consider the seamless integration of all resources from the local to the international, from academic to commercial, from single PI to multi-national collaborations, from life sciences to physical sciences, to be a major CI opportunity and challenge to accelerate open science.

CI here includes hardware resources, data resources, reproducible workflows, software, and human expertise. To achieve this integration requires a set of collaborative national-scale projects that involve scientists, research computing facilitators, software developers, and IT infrastructure professionals.

The distributed high-throughput computing (dHTC) ecosystem created by the Open Science Grid (OSG) provides a proof of principle that large-scale resource sharing enables science. In the OSG Ecosystem Big Science facilities like the Large Hadron Collider and LIGO play a special role in driving the CI, while the domain science independence of OSG guarantees its openness to all of science.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

The open intellectual sharing implicit in multidisciplinary, multi-institutional collaborations has led to some of the greatest scientific breakthroughs of modern science. A CI ecosystem that is nimble and flexible enough to federate voluntarily contributed resources is required to complement the intellectual sharing and collaborative relationships that are essential to multidisciplinary science.

Open science often exhibits enormous data-intensive workflows that may be decomposed into a huge number of individually manageable components. These workflows perform analyses utilizing enormous shared datasets, often through multiple filters in a pipeline, to uncover scientific insights. Current examples include a broad variety of biological investigations that require ensembles of databases such as those obtainable via the National Center for Biotechnology Information (NCBI). Other projects like the Large Hadron Collider (LHC), the LIGO Scientific Collaboration, Ice Cube, Alpha Magnetic Spectrometer (AMS), the Large Synoptic Sky Survey, the XENON1T experiment and its NT follow-up, and the Stage-4 ground-based cosmic microwave background (CMB-S4) program of telescopes as well as many others each have their own high impact science goals that the proponents of these experiments are much better positioned to defend than we are. The purpose of this letter is not to explain the virtue of any particular science or engineering challenge, but rather to point out their technical interdependencies and commonalities as far as Cyberinfrastructure (CI) is concerned.

All of the aforementioned science projects and many others continue to benefit greatly from advances of the state of the art in distributed High Throughput Computing (dHTC). Larger experimental programs typically require larger investments into software R&D in particular, and CI in general. Within the Open Science Grid (OSG) Consortium, we have demonstrated that not only can Big Science projects benefit from each other's CI developments, but those very same developments are also beneficial to midscale science projects and ultimately the long tail of science across research campuses at all scales. The global computing infrastructure created for the LHC is now being used by the LIGO Scientific Collaboration, Ice Cube, and other scientific or campus communities. Universities and National Laboratories worldwide that collaborate with these US lead science programs share their local computing resources via the interfaces developed and maintained for the LHC program. Those not already participating in the LHC are deploying OSG interfaces to integrate their local resources into their global collaborations. Experiments like XENON1T that are located in Europe or elsewhere use the LHC CI developments to bring their data to the US, and analyze it across the shared US infrastructure originally developed, and still dominated in its use, by the LHC.

XENON1T in particular archives its data at an LHC Tier-1 center in Europe because that center has an institutional commitment to XENON1T. The OSG user support organization at University of Chicago collaborates with XENON1T to operate an instance of Rucio, the ATLAS data management software, developed largely at CERN, to bring XENON1T data to the US. From there, XENON1T processing is orchestrated using the glideinWMS infrastructure developed for CMS and jointly operated by OSG and CMS. XENON1T uses a combination of an XSEDE allocation, U. Chicago campus resources, and resources shared by others on OSG. The XSEDE allocation was applied for by Wuerthwein in the name of OSG, in support of a new junior faculty in the UCSD physics department working on XENON1T. This is just one of many examples of CI reuse within the OSG Consortium. We picked this example because XENON1T is a midscale experiment that would never have had the funding to create the mechanisms they use by themselves, and because it shows the kind of incentive structures crucial for open science collaborations across institutions. Institutional commitments to both the science and the open consortium led to the collaborative effort. Without the open consortium, XENON1T's imminent first publication would have had to have been done very differently.

Similarly, LIGO's operations via the OSG infrastructure start by importing their data into the University of Nebraska-Lincoln (UNL) 100Gbps connected data infrastructure. UNL does not have any faculty members affiliated with LIGO. Its connection to LIGO is entirely via the OSG Consortium. From UNL, pegasus, software that helps workflow-based applications execute in a number of different highly distributed environments, pulls the data to clusters operated by OSG worldwide and directs its processing using a technology called HTCondor. These processors are provisioned by the same glideinWMS system mentioned previously. In many cases network routing established via LHCONE is used to do so because both endpoints are part of the global computing infrastructure for the LHC.

These successes in cross cutting use of CI were made possible by the open science vision of a collaboration of computer scientists, IT infrastructure professionals, and scientists within the LHC community more than 10 years ago. They insisted to not just build something that serves the LHC, but to embed it into an open consortium that serves all of open science. The successes of the last 10 years were made possible because the NSF shared this vision that transcends narrow science silos and budgets spent entirely according to "mission mandates".

We are concerned that the NSF has lost this vision. We are concerned that new programs are being designed much more narrowly, and without any mandate to serve all of science. In fact, we have been told by our program managers that our vision of cross cutting CI for open science is suffering the tragedy of the commons. Individual areas within NSF are each focused on their own narrow objectives at a time when university administrations nationwide are recognizing the need for CI that is common to all of science. Now is the time to aggressively advance integration of CI investments nationwide, rather than retrench into narrowly defined silos.

An effective means to transparently combine cyberinfrastructure (CI) from multiple collaborating resource providers remains a key need. Scientists should be able to submit locally, and compute globally. Integration is needed for many of the components of modern CI including both computing and data resources. This is particularly important as universities increasingly view CI as a common-good resource, analogous to (and increasingly augmenting) the library. There are currently few to no options that allow researchers to develop, deploy and scale workflows via a scalable resource which incorporates a CI component over which they maintain control, ranging from a local desktop to a campus cluster or data store. Commercial clouds meet some of these goals, but they lack the ability to incorporate locally deployed resources into the same workflow; at present, they are also not always cost effective. dHTC is a technology that can meet the above needs economically and effectively.

NSF is the federal agency with the broadest research mission across fields of fundamental science and engineering. NSF is moreover committed to research being fully integrated with education and training of tomorrow’s top scientists and engineers. This mission and commitment positions the NSF to be the primary driver of an Open Science CI ecosystem that integrates resources across campuses of all scales to solve science problems at all scales, from the largest Big Science projects like the High Luminosity LHC to single PIs at liberal arts colleges.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

Technical challenges confronting the creation of an open, federated Open Science CI ecosystem are considerable, but it’s equally necessary to build a coalition of the willing that serves key stakeholders and is accessible to the broad range of scientists conducting data intensive research today.

Technical Challenges:

The overarching goal is to create end-to-end solutions that decrease the time to science. This is composed of technical challenges to be delineated below, with the emphasis that at the end of the day, the greatest need is for reliable services that help scientists to affordably do science while not being required to navigate between multiple disparate resources, or worse, to develop their own CI. The final integrated distributed system will maintain local ownership and autonomy while facilitating collaborative science.

Access should be made as easy as possible, with complexity only added as needed. Here solutions such as InCommon are possible ways to allow broad participation.
Enabling campus researchers to share computational and data resources with external collaborators is a powerful multiplier in advancing science. Sharing spare capacity for even short durations allows an institutional HPC resource a cost-efficient means of participating in a larger cyber ecosystem. OSG allows institutions to integrate their HPC cluster resource to support collaborative, multi-institutional science.

Integration is not only needed between different institutions but also between different layers of the CI -- in particular, networking, data and processing. Data handling is even more challenging than processing. End to end performance depends upon all of the three above components, and personnel with systems expertise are required to provide a transparently scalable, fault-tolerant resource between sites. The OSG Consortium has made considerable progress toward this goal in cooperation with international scale big science projects.

The data layer will need to incorporate multiple mechanisms to deal with the differing demands of various workflows. Caching, streaming and transferring data will need to be utilized to implement federated data ecosystems that support multidisciplinary science. Tools such as Globus Connect, GridFTP, http squid caches and xrootd will put stress on the network. This in turn will make significant demands upon the network, leading to use of PerfSONAR or software defined networks to ensure performance. Good, portable software is needed; singularity is a means toward this end.

Organization and Sustainability:

The vast majority of CIOs and campus decision-makers nationwide have now understood that Research Computing is a common good similar to the library in that it must be coherently organized and sustainably funded for the research mission of the campus to thrive in the digital age. A very large investment is thus being made into campus CI nationwide that in aggregate exceeds easily what the NSF could possibly sustain. In fact, it is not uncommon today for large research universities to operate compute and storage infrastructures that are as large or larger than the NSF's Track 2 systems that are part of the national XRAC allocation mechanism.

The next step is for the NSF to recognize that science is a team sport that very often crosses institutional in addition to disciplinary boundaries. The investments that individual campuses have been making require national integration in order to fully realize their potential scientific impact. The NSF is uniquely positioned to achieve such integration. Other agencies have missions that are too narrowly focused to even aspire to the vision that must be executed. It is not sufficient to focus only on integrating the CI investments made by the NSF itself. To the contrary, all CI investments that serve open science ought to be within scope for such an integration, regardless of their source of funding.

To achieve this vision requires a mix of unfunded consortia, i.e. collaboration of the willing under a common vision, and NSF funded projects that can implement the vision. Examples of such relationships between consortium and funded project are the Campus Research Computing (CaRC) consortium and its relationship to ACI-REF, or the Open Science Grid (OSG) Consortium and its relationship to the OSG funded project. OSG is a functioning federation with a decade of experience in running a shared, national infrastructure.

An excellent further example of a broad coalition serving multiple stakeholders is the LHC community. For decades this international community has managed to cooperatively fund and manage some of the largest, most complex CI environments in production today. In the US, cooperation between national laboratories and universities has been forged and maintained, leading to historic scientific breakthroughs. OSG has leveraged several advances that have originated in the LHC experiments to broadly impact other unrelated areas of science. This serves as a strong indicator that an open, federated CI can concurrently serve numerous areas of science.

Social Challenges:

We'd be remiss if we didn't acknowledge that there are social challenges to overcome while creating such an integrated infrastructure. It's a fact that not all people learned in kindergarten that sharing is a good thing. Some need to be educated later in life, and the NSF has a role to play in that education, e.g. by encouraging resource sharing as part of solicitations. To make sharing work, we learned that some basic concepts are essential. First among them is the notion that “owners rule”. Any owner must be allowed to take away any resource at any time. This implies that resource use policies must be configurable locally, and preemption requires fault tolerance at the workflow and dataflow execution layer. Overcoming social challenges thus leads to technical requirements on the CI. Last year OSG facilitated the voluntary sharing of 176M cpu hours from locally owned resources. All of these resources, most of them from campuses, were shared in accord with the philosophy described here. The capacity of other untapped campus resources, which no doubt far exceeds this number, could represent a tremendous national resource.

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

Training in, and the sharing of, best practices and workforce development is a natural outcome of shared CI resources that should be leveraged. Tuning the performance of a distributed infrastructure stretches technical expertise and grows an integrated workforce which necessarily must interact and jointly solve problems. Several national organizations such as ACI-REF, CaRC, CASC, EDUCAUSE and XSEDE illustrate the professional benefits of shared expertise. These benefits extend to engineers developing and maintaining shared distributed hardware and software systems. OSG holds periodic training sessions for personnel from multiple institutions, who gather to acquire the skills necessary to deploy and maintain systems serving a wide variety of scientists. Further training has enlisted the ACI-REFs, training them to assist scientists to understand and utilize the current production dHTC (distributed High Throughput Computing) environment. This has broadened awareness of the power and availability of dHTC, especially via OSG.

Training is also needed to assist data-intensive scientists in their work. Several approaches have been used to assist researchers whose codes and workflows may be partitioned, creating a distributed load that maps effectively to the CI maintained by OSG. Weeklong summer schools are held annually, drawing students and staff from institutions distributed internationally. Software Carpentry Workshops, an international training effort gaining popularity worldwide, have been extended to include training in OSG and Globus Connect. Once researchers are introduced to the power of dHTC, novel workflows often are devised, resulting in new demands from the existing CI. This in turn further challenges the systems developers of the CI itself, and iterative development continues.

Consent Statement